

Access Control Policies

These topics describe how to work with access control policies:

- Access control policies, on page 1
- Requirements and prerequisites for access control policies, on page 6
- Manage access control policies, on page 6

Access control policies

An access control policy is a hierarchical policy that allows you to specify, inspect, and log network traffic. You can use the characteristics of each connection to allow, trust, block, or monitor the connection.

Each managed device can be assigned to one access control policy. The data that the policy's assigned (also known as targeted) devices collect about your network traffic can be used to filter and control that traffic based on:

- simple, easily determined transport and network layer characteristics: source and destination, port, protocol, and so on
- the latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application used, or URL visited
- realm, user, user group, or ISE attribute
- custom Security Group Tag (SGT)
- characteristics of encrypted traffic; you can also decrypt this traffic for further analysis
- whether unencrypted or decrypted traffic contains a prohibited file, detected malware, or intrusion attempt
- time and day (on supported devices)

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance. For example, reputation-based blocking uses simple source and destination data, so it can block prohibited traffic early in the process. In contrast, detecting and blocking intrusions and exploits is a last-line of defense.

Access control policy components

This list describes the main components of an access control policy.

- Name and description: Each access control policy must have a unique name. A description is optional.
- Inheritance settings: Policy inheritance allows you to create a hierarchy of access control policies. A parent (or base) policy defines and enforces default settings for its descendants. A policy's inheritance settings allow you to select its base policy. You can also lock settings in the current policy to force any descendants to inherit them. Descendant policies can override unlocked settings.
- Policy assignment: Every access control policy identifies which devices use it. Each device can use only one access control policy. You can also assign the policy to device templates.
- Rules: Access control rules provide a granular method of handling network traffic. Rules are numbered, starting at 1, including rules inherited from ancestor policies. The system matches traffic to rules from the top down, ascending by rule number.

  Traffic is handled by the first rule for which all conditions match. Conditions can be simple or complex, and often depend on licensing.
- Default action: The default action determines how the system handles and logs traffic that is not handled by any other access control configuration. It can block or trust all traffic without further inspection, or inspect traffic for intrusions and discovery data. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.
- Security Intelligence: This feature acts as the first line of defense against malicious internet content. It enables blocking of connections based on IP address, URL, and domain name reputation intelligence. To ensure continual access to vital resources, you can override block list entries with custom do not block list entries.
- HTTP responses: When a user's website request is blocked, you can display either a generic response page or a custom page. Alternatively, you can warn users but allow them to proceed to the requested site.
- Logging: Logging settings enable configuration of default syslog destinations for the policy. These settings apply to the policy and all included decryption, prefilter, and intrusion policies unless overridden in individual rules or policies.
- Advanced access control options: Advanced policy settings typically require little modification, and defaults are often suitable. Modifiable options include traffic preprocessing, decryption, identity, and various performance configurations.

Access control policy default actions

The default action is the access control policy configuration that determines how the system handles and logs traffic that is not handled by any other access control configuration. The default action can block or trust all traffic without further inspection, or inspect traffic for intrusions and discovery data.

In a simple access control policy, the default action specifies how a device handles all traffic. In a more complex policy, the default action handles traffic that:

- is not fast pathed by the prefilter policy
- is not on a Security Intelligence block list
- is not blocked by the decryption policy (encrypted traffic only)
- matches none of the rules in the policy (except monitor rules, which match and log—but do not handle or inspect—traffic)

Note You cannot perform file or malware inspection on traffic handled by the default action. Logging of connections handled by the default action is initially disabled but can be enabled.

If you are using policy inheritance, the default action for the lowest-level descendant determines how traffic is handled. An access control policy can inherit its default action from the base policy, but this inheritance is not enforced.

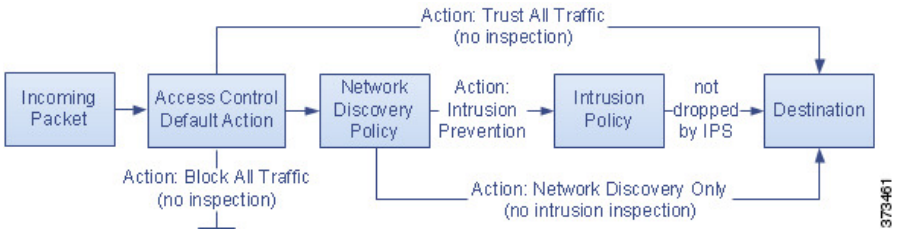
This table describes the types of inspection you can perform on traffic handled by each default action.

Table 1: Types of default actions and inspection options

Default action	Effect on traffic	Inspection type and policy
Access Control: Block All Traffic	Block without further inspection.	None.
Access Control: Trust All Traffic	Trust (allow to its final destination without further inspection).	None.
Intrusion Prevention	Allow, as long as it is passed by the intrusion policy you specify.	Intrusion, using the specified intrusion policy and associated variable set. Discovery, using the network discovery policy.
Network Discovery Only	Allow.	Discovery only, using the network discovery policy.
Inherit from base policy	Defined in base policy.	Defined in base policy.

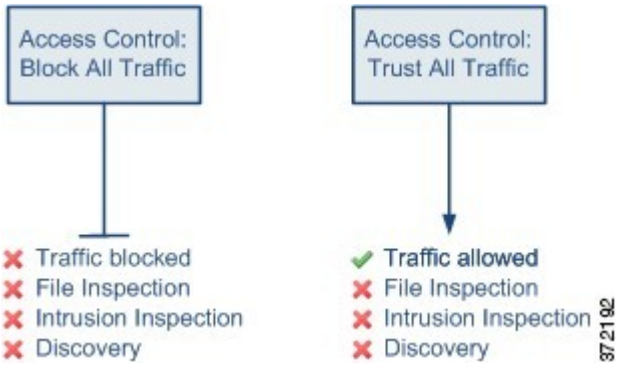
This diagram illustrates the table.

Figure 1: Default action configuration options



This diagram illustrates the **Block All Traffic** and **Trust All Traffic** default actions.

Figure 2: Block all traffic and trust all traffic default actions



This diagram illustrates the **Intrusion Prevention** and **Network Discovery Only** default actions.

Figure 3: Intrusion prevention and network discovery only default actions



Note The purpose of **Network Discovery Only** is to improve performance in a discovery-only deployment. Different configurations can disable discovery if you are only interested in intrusion detection and prevention.

Access control policy inheritance

Access control policy inheritance is a hierarchical policy-based implementation that allows child policies to inherit rules and settings from ancestor or base policies. This structure ensures that mandatory rules defined in parent policies are consistently enforced, while providing the flexibility for child policies to selectively override default configurations.

Access control policy hierarchy structure

You can nest access control policies, where each policy inherits the rules and settings from an ancestor policy. You can enforce this inheritance, or allow lower-level policies to override their ancestors.

Just as you create a domain hierarchy, you can create a corresponding hierarchy of access control policies. A descendant, or child, access control policy inherits rules and settings from its direct parent, or base, policy. That base policy may have its own parent policy from which it inherits rules and settings, and so on.

An access control policy's rules are nested between its parent policy's Mandatory and Default rule sections. This implementation enforces Mandatory rules from ancestor policies, but allows the current policy to write rules that preempt Default rules from ancestor policies.

You can lock these settings to enforce them in all descendant policies:

- Security Intelligence—Connections that are allowed or blocked based on the latest reputation intelligence for IP addresses, URLs, and domain names.
- HTTP response pages—Displaying a custom or system-provided response page when you block a user's website request.
- Advanced settings—Specifying associated subpolicies, network analysis settings, performance settings, and other general options.

When using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

Rule and other policy warnings

Policy and rule editors use icons to mark configurations that could adversely affect traffic analysis and flow. Depending on the issue, the system may warn you when you deploy or prevent you from deploying entirely.

Note Move your pointer over an icon to read the warning, error, or informational text.

Table 2: Policy error icons

Icon	Description	Example
Errors ( )	If a rule or configuration has an error, you cannot deploy until you correct the issue, even if you disable any affected rules.	A rule that performs category and reputation-based URL filtering is valid until you assign a device that does not have a URL Filtering license. At that point, an error icon appears next to the rule, and you cannot deploy until you edit or delete the rule, unassign the device, or enable the license.
Warning ( )	You can deploy a policy that displays rule or other warnings. However, misconfigurations marked with warnings have no effect. If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue.	Preempted rules or rules that cannot match traffic due to misconfiguration have no effect. This includes conditions using empty object groups, application filters that match no applications, excluded LDAP users, invalid ports, and so on. However, if a warning icon marks a licensing error or model mismatch, you cannot deploy until you correct the issue.

Icon	Description	Example
Information ()	Information icons convey helpful information about configurations that may affect the flow of traffic. These issues do not prevent you from deploying.	The system might skip matching the first few packets of a connection against some rules, until the system identifies the application or web traffic in that connection. This allows connections to be established so that applications and HTTP requests can be identified.
Rule Conflict ()	When you enable rule conflict analysis, this icon appears in the rule table for rules that have conflicts.	Conflicts include redundant rules, redundant objects, and shadowed rules. Redundant and shadowed rules do not match traffic because previous rules would already match the criteria. Redundant objects make your rules unnecessarily complex.

Requirements and prerequisites for access control policies

Ensure that your environment meets these requirements before configuring access control policies.

- You must have these user privileges to configure and deploy access control policies:
  - **Admin**, **Super Admin**, and **Edit Only** privilages allow you to modify access control policies.
  - **Admin**, **Super Admin**, and **Deploy Only** privilages allow you to deploy queued access control policy changes.
  - If you have **Edit Only** privilages, you need deployment permissions to deploy your changes.

Manage access control policies

Configure and manage access control policies by editing system-provided policies and creating custom access control policies.

Procedure

Step 1 Choose **Policies** > **Security policies** > **Access Control**.

At the top of the page, there are convenient links to some related features: Object management, Intrusion policies, Network Analysis policies, DNS policies, and policy Import/Export.

Step 2 These actions help you to manage your policy:

- Create—Click **Create Policy**; see Creating a basic access control policy, on page 7.
- Edit—Click **Edit** (); see Editing an access control policy, on page 8
- Delete—Click **Delete** (). You must remove any device assignments before deleting a policy.

To delete more than one policy at a time, select the check boxes for the policies, then select **Delete Policies** above the table.

- Copy—Select **Clone** from the **More** (⋮) menu. Device assignments are not retained in the copy.
- Lock or unlock a policy—See Lock an access control policy, on page 10.

Step 3 These options help you to analyze, monitor, and customize your policies:

- Analyze policy—Select one or more policy and then click **Analyze Policies** to evaluate access control policies for anomalies such as redundant or shadowed rules, and take action to fix discovered anomalies. The analysis job is sent to the cloud and takes time to complete. See Identify and fix anomalies with policy analyzer & optimizer, on page 20.

  The Anomalies column shows the results of the analysis. Click the **% Optimizable** link to see the anomalies, or **Re-Analyze** to run the analysis again. The Last Analyzed column shows when Policy Analyzer & Optimizer was last run.

  After completing analysis and optimization, you can download the reports by selecting the following options from the **More** (⋮) menu: **Download Last Policy Analysis**, **Remediation History**.

Note To use the policy analysis feature, you must be using a Cloud-Delivered Firewall Management Center or an On-Prem Firewall Management Center connected to Security Cloud Control (Security Cloud Control). If your setup does not meet requirements, the explanatory dialog that opens when you click this button includes an **Integrate** button to help get you started. Policy Analyzer & Optimizer operates in the cloud only.

- Columns—Click the **Show/Hide Columns** icon above the list of rules to select which information to show in the table. Click **Show All/Hide All** to quickly add or remove all listed columns, excepting name and actions. Click **Default** to undo all of your customizations.
- Inheritance—Vertical lines indicate inheritance relationships between policies.
- Report—Select **Generate Report** from the **More** (⋮) menu. The report is generated as a background process. Go to the message/notifications center and look in the tasks list. When the report is finished, you can download it from the notification.
- View the audit log—Click **View Audit Log** from the **More** (⋮).

Creating a basic access control policy

When you create a new access control policy, it contains default actions and settings. After creating the policy, you are immediately placed in an edit session so that you can adjust the policy to suit your requirements.

Procedure

Step 1 Choose **Policies** > **Security policies** > **Access Control**.

Step 2 Click **Create Policy**.

Step 3 Enter a unique **Name** and, optionally, a **Description**.

Accented and other non-ASCII characters are not recommended in access control policy names.

Step 4 Optionally, choose a base policy.

If you select a base policy, the base policy defines the default action and you cannot select a new one in this dialog box. Logging for connections handled by the default action depends on the base policy.

Step 5 When you do not select a base policy, specify the initial **Default Action**:

- **Block all traffic** creates a policy with the **Access Control: Block All Traffic** default action.
- **Intrusion Prevention** creates a policy with the **Intrusion Prevention: Balanced Security and Connectivity** default action, associated with the default intrusion variable set.
- **Network Discovery** creates a policy with the **Network Discovery Only** default action.

When you select a default action, logging of connections handled by the default action is initially disabled. You can enable it later when you edit the policy.



Note If you want to trust all traffic by default, or if you chose a base policy and do not want to inherit the default action, you can change the default action later.

Step 6 Optionally, choose the devices to assign to the policy. To narrow the devices that appear, enter a search string. The list includes both devices and device templates.

If you want to deploy this policy immediately, you must perform this step.

Step 7 Click **Save**.

The new policy opens for edit. You can add rules to it and make other changes as needed. See Editing an access control policy, on page 8.

Editing an access control policy

When you edit an access control policy, you should lock it to ensure that your changes do not get overridden by another person who might edit it simultaneously.

To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Note When you edit a policy and do not lock it, a banner message might indicate that other users are currently editing the policy. When one of these users saves changes, you will be prompted to either merge or discard your changes. You should immediately decide what to do. For more information, see Merge concurrent edits to an access control policy, on page 11 .

Procedure

Step 1 Choose **Policies** > **Security policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

Step 3 Edit your access control policy.

Note You can operate on multiple rules at one time by selecting their checkboxes in the left column, then selecting the action you want to perform from the **Select Bulk Rule Actions** drop-down list next to the search box. Bulk editing is available for enabling and disabling, copying, moving, deleting, and editing rules, or viewing hit counts or related events. You can also remove object overlaps in the selected rules.

You can change the following settings or perform these actions:

- Name and description—Click **Edit** (✎) next to the name, make your changes, and click **Save** .
- Default action and settings—Choose a value from the **Default Action** drop-down list, then click **Cog** (⚙) , make your changes to the settings, and click **OK** . For detailed information, see Configure the access control default action, on page 13 .
- Associated policies—To edit or change policies in the packet flow, click the policy type in the packet flow representation below the policy name. You can select the **Prefilter Rules** , **Decryption** , **Security Intelligence** , and **Identity** policies. When necessary, click **Access Control** to return to the access control rules.
- Policy assignment—To identify the managed devices targeted by this policy, click the **Assigned: x devices** link. You can assign the policy to devices or device templates.
- Rules—To manage access control rules, and to inspect and block malicious traffic using intrusion and file policies, click **Add Rule** , or right-click an existing rule and select **Edit** or another appropriate action. The actions are also available from the **More** (⋮) button for each rule. See Create and edit access control rules .
- Layout—Use the **Grid/Table View** icon above the list of rules to change the layout. Grid view provides color-coded objects in an easy-to-see layout. Table view provides a summary list so that you can see more rules at once. You can freely switch views without impacting the rules.
- Columns (table view only)—Click the **Show/Hide Columns** icon above the list of rules to select which information to show in the table. Click **Show/Hide Empty Columns** to quickly add, or remove, all columns that have no information, that is, you are not using those conditions in any rule. Click **Revert to Default** to undo all of your customizations.

- Analyze rule logic—You can select the following options from the **Analyze** menu to examine the logic of your rules:
  - **Manage Rule Hit Counts** —To view statistics on how many connections matched each rule. See View rule hit counts, on page 21 .
  - **Enable/Disable Rule Conflicts** —Select whether you want to see information on whether rules interfere with each other. You can then view the results using the following commands. See Analyzing rule conflicts and warnings, on page 23
    - **Show Warnings and Errors** —See whether there are rules with configuration issues that you need to address.
    - **Show Policy Warnings** —See whether there are configuration issues with the policy.
    - **Show Rule Conflicts** —See whether you have redundant or shadowed rules. These conflicts could prevent certain rules from ever being matched by connections, meaning either that you need to fix the match criteria, move the rule, or simply delete the rule.

Step 4 To change additional settings for the policy, select one of these options from the **More** drop-down arrow at the end of the packet flow line.

- **Encrypted Visibility Engine** —Whether to enable and configure the encrypted visibility engine. See Configure EVE .
- **Advanced Settings** —To set preprocessing, decryption, identity, performance, and other advanced options. See Configure advanced settings for the access control policy .
- **HTTP Responses** —To specify what the user sees in a browser when the system blocks a website request. See Choose HTTP response pages .
- **Inheritance Settings** —To change the base access control policy for this policy, and to enforce this policy's settings in its descendant policies. See Choose a base access control policy, on page 15 and Lock settings in descendant access control policies, on page 16 .
- **Logging** —To set the default logging options for the policy.

Step 5 Click **Save** .

Lock an access control policy

Lock an access control policy to prevent other administrators from editing it, ensuring that your changes will not be invalidated if another administrator edits the policy and saves changes before you save your changes.

If you do not lock a policy, you might have to merge your changes if another user edits the policy simultaneously and saves changes first. For more information, see Merge concurrent edits to an access control policy, on page 11.

The lock is for the access control policy itself. The lock does not apply to objects used in the policy. For example, another user can edit a network object that is used in a locked access control policy. Your lock remains in place until you explicitly unlock the policy, so you can log out and come back to your edits later.

When locked, other administrators have read-only access to the policy. However, other administrators can assign a locked policy to a managed device.

Before you begin

Any user role that has permission to modify the access control policy has permission to lock it, and to unlock a policy that was locked by another user.

However, the ability to unlock a policy that was locked by another administrator is controlled by the following permission: **Policies** > **Security policies** > **Access Control** and choose the **Access Control Policy** > **Modify Access Control Policy** > **Override Access Control Policy Lock**.

Procedure

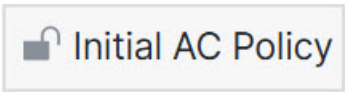
Step 1 Choose **Policies** > **Security policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to lock or unlock.

In the access control policy list:

- A lock icon next to a policy name indicates that the policy is locked. Hover over the icon to see which user locked the policy.

Step 3 Click the lock icon next to the policy name to lock or unlock the policy.



If the policy inherits settings from a parent policy, you must choose one of these options when you click the lock icon.

- **Lock/Unlock this policy**—The locking or unlocking is for this policy only.
- **Lock/Unlock all policies in hierarchy**—This policy and all parent policies are locked or unlocked. If a parent policy is already locked by another administrator, you will see a message and you will not be able to lock that parent policy. When unlocking policies, if you have the Override Access Control Policy Lock permission, all parent policies are unlocked even if they were locked by other users.

Merge concurrent edits to an access control policy

Handle conflicts when multiple users make changes to the same access control policy simultaneously, ensuring no critical changes are lost while maintaining policy integrity.

When you allow multiple users to edit an access control policy, you have two choices:

- Have users lock the policy when editing it. This ensures that a single user at a time can make changes. Using locking, you never have to worry about lost or conflicting changes. However, it limits your ability to respond quickly to unexpected events, especially if a user leaves the policy locked while out of the office. For more information about locking policies, see Lock an access control policy, on page 10.

- Allow users to simultaneously edit the policy. If more than one user has made unsaved edits at a given time, the first user who saves the policy gets all of their changes saved. Other users are immediately notified of the save, and must then merge their changes. The rest of this topic explains this approach.

Note If you are using change management, where edits are done within the scope of a ticket, the ticket locks the policy and concurrent editing is not possible.

Before you begin

To limit the amount of merging you have to do, save the policy early and often.

Procedure

Step 1 Determine if other users are also editing the policy.

When you edit an access control policy, look for this message banner at the top of the page:

```
User name is currently editing this policy.
```

This message indicates that the named user (there might be more than one) currently have the policy open for edit.

Step 2 Watch for notifications that another user saved changes and take action.

This banner message indicates another user saved changes and you need to take immediate action:

```
User name modified this policy and saved changes. Merge Discard.
```

Click one of these links:

- **Merge**—Open a merge window where you can make decisions about which of your changes to keep or to discard, and to identify changes that cannot be merged. See the next step when taking this option.
- **Discard**—Discard all of your changes immediately and start over. If you select this option, the page is refreshed with the latest changes from the other user.

Step 3 Merge changes.

When you click the **Merge** option, a merge window opens with a summary of the total number of observed differences between your edits and the last user's saved edits. No other user's unsaved edits are included. Observations include direct conflicts and informational notifications.

Note You must complete all merge decisions before you close the window. You cannot wait until later once you click the **Merge** option.

a) Evaluate each conflict and make a decision for each.

When evaluating the change list, consider these points:

- Initially, all observations are shown, but you can select/deselect the filtering options for **Conflict** and **Info** to limit the displayed information. Conflicts arise when users edit the same rule or setting. Informational items are for edits to different rules or settings.
- If there are no conflicts, your cache is immediately updated with the last saved changes and you can proceed. For example, if you are editing rule 1, and the other user saved changes to rule 2, it is unlikely that there are any conflicts.
- For each conflict (other than rule name), the window shows the version of the element saved by the other user (**Version on Firewall Management Center**), and your changes (**Modified Version**). Changes are color-coded for new (e.g. an option that was not previously defined), edited, or deleted policy elements. You must select either to **Discard** your change, or to **Accept Mine** and overwrite the other user's change. When discarding a change, the rule is refreshed with the other user's saved change and yours is removed.
- Rule conflicts are considered at the rule level, not per element within the rule. For example, if you edit destination networks, and the saved changes are for source networks, either your changes or the other user's changes are retained, not both.
- If the conflict involves the name of a rule, the system tries to create a unique name by adding an underscore and number to the name, such as Rule_1. You can alternatively enter a new name in the edit box provided and click **Save**. If you click **Discard**, the generated rule name is used.
- You cannot skip any conflict. For each conflict, you must make a decision to discard or keep your change.
- If the user who saved changes deleted a rule that you edited, the rule is deleted and you do not have the option to keep the rule.
- If the user who saved changes updated a rule, and your change is to delete the rule, the rule is retained: you cannot merge your deletion.
- For these options, if your change conflicts with the saved changes, your change is discarded and merging the change is not an option:
  - Other policy assignments: prefilter, Security Intelligence, identity, decryption
  - Changes to advanced settings, policy logging, EVE, HTTP response pages.
  - Policy default action.
  - Inheritance settings.

b) Click **Close**.

Step 4 Once the merger is complete, you can either save the policy immediately or continue editing it. Saving the policy might reduce the impact of future merges.

Configure the access control default action

Configure the default action for an access control policy when connections do not match specific access control rules.

The default action for an access control policy is applied to any connection that:

- is not fast pathed by the prefilter policy
- is not on a Security Intelligence block list
- is not blocked by the decryption policy (encrypted traffic only)
- matches none of the rules in the policy (except monitor rules, which match and log—but do not handle or inspect—traffic)

Procedure

Step 1 Edit the access control policy whose inheritance settings you want to change; see Editing an access control policy, on page 8.

Step 2 Select the **Default Action** at the bottom of the rules list.

For detailed information on what each option does, see Access control policy default actions, on page 2.

Step 3 Click **Cog** (⚙) to configure the default action settings.

You can configure the following options. Click **OK** when finished.

- Logging options—Whether to log the connection. You can **Log at beginning of connection**, **Log at end of connection**, or both. If you select block as the default action, you can log at the beginning of the connection only.
- **Send connection events to**—If you select one of the logging options, select whether to send events to any combination of the following:
  - **Event Viewer**—Send events to Cisco Cloud and view them in the event pages in Security Cloud Control.

Note To view events in Security Cloud Control, you must enable policy logging and configure the **Cisco Cloud Events** setting in Security Cloud Control for your device and event types. Enabling logging in an access control rule or default action alone does not configure the devices to send events to Cisco cloud. For more information about configuring **Cisco Cloud Events** setting, refer to Send Firewall Threat Defense event logs to Security Cloud Control using a direct connection.

  - **Syslog server**—Send events to the default syslog server configured for the policy. You can configure overrides to specify a different severity level or syslog server destination.

  (Splunk or SIEM syslog server integration) If you have configured Cloud-Delivered Firewall Management Center as the source of connection events in Splunk configuration, choose **Firewall Management Center** under **Send Connection Events to** options. If you have configured Firewall Threat Defense as the connection event source, select at least one syslog destination. For Splunk or any SIEM syslog configuration procedure, see the *Cisco Secure Firewall Management Center Administration Guide*.

- **SNMP trap**—If you enable logging, you can send SNMP traps to an SNMP server. Select an SNMP configuration, or click + to configure a new one.

- **Default action variable set**—If you selected one of the intrusion prevention default actions, select the variable set that should be used with the intrusion policy you selected.

Step 4 Click **Save**.

Manage access control policy inheritance

Inheritance relates to using another policy as a base policy for an access control policy. This allows you to use one policy to define some baseline characteristics that can be applied to multiple policies. To understand how inheritance works, refer to Access control policy inheritance, on page 4.

Procedure

Step 1 Edit the access control policy whose inheritance settings you want to change; see Editing an access control policy, on page 8.

Step 2 Manage policy inheritance:

- Change base policy — To change the base access control policy for this policy, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line and proceed as described in Choose a base access control policy, on page 15.
- Lock settings in descendants — To enforce this policy's settings in its descendant policies, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line and proceed as described in Lock settings in descendant access control policies, on page 16.
- Inherit settings from base policy — To inherit settings from a base access control policy, proceed as directed in Inherit access control policy settings from the base policy, on page 16.

Choose a base access control policy

Choose a base access control policy to establish policy inheritance, allowing a child policy to inherit rules and settings from its parent while maintaining the ability to modify unlocked configurations.

When you change the base policy for the current access control policy, the system updates the current policy with any locked settings from the new base policy.

Procedure

Step 1 In the access control policy editor, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 Choose a policy from the **Select Base Policy** drop-down list.

To remove inheritance, select **None**.

Step 3 Click **OK**.

Step 4 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Lock settings in descendant access control policies

Lock a setting in an access control policy to enforce the setting in all descendant policies.

Descendant policies can override unlocked settings. When you lock settings, the system saves overrides already made in descendant policies so that the overrides can be restored if you unlock settings again.

Procedure

Step 1 In the access control policy editor, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 In the Child Policy Inheritance Settings area, check the settings you want to lock.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

Step 3 Click **OK** to save the inheritance settings.

Step 4 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Inherit access control policy settings from the base policy

Configure a child access control policy to inherit settings from a base policy or override them with custom configurations.

Child access control policies inherit many settings from their base policy by default. If these settings are unlocked in the base policy, you can override them.

If you later reinherit the settings from the base policy, the Cloud-Delivered Firewall Management Center displays the base policy's settings and dims the controls. However, it saves the overrides you made and restores them if you disable inheritance again.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**, or select **HTTP Responses**, **Logging**, **Encrypted Visibility Engine**, or **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line

Step 2 Check the **Inherit from (*base policy*)** check box for each setting you want to inherit.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

Step 3 Click **Save**.

What to do next

- Deploy configuration changes.

Assign access control policy to devices

An access control policy specifies the devices that use it. Each device can be assigned to only one access control policy. You can also assign the policy to device templates. Templates are included in the list of available and selected devices.

Procedure

Step 1 In the access control policy editor, click the **Assigned to x devices** link.

Step 2 Build your target list:

- Add — Select one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**.
- Delete — Click **Delete** ( ) next to a single device, or select multiple devices, right-click, then choose **Delete Selection**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Under **Impacted Devices**, the system lists the devices whose assigned access control policies are children of the current policy. Any change to the current policy affects these devices.

Step 3 Assign the access control policy to devices:

- Add — Select one or more **available devices**, then click the **Right Arrow** ( ) icon or drag and drop into the list of **assigned devices**.
- Delete — Click the **Delete** ( ) next to a device, or select multiple devices, and then click the Left Arrow .
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Step 4 Click **OKApply** to save your assigned device settings.

Step 5 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Logging settings for access control policies

Logging settings for access control policies are configuration options that allow you to configure default syslog destinations for the current access control policy. The settings are applicable to the access control policy and all the included decryption, prefilter, and intrusion policies unless the syslog destination settings are explicitly overridden with custom settings in included rules and policies.

To configure logging settings for an access control policy, select **Logging** from the **More** drop-down arrow at the end of the packet flow line.

Logging for connections handled by the default action is initially disabled.

Intrusion settings and file and malware settings are effective only after you have selected an

option at the top of the page for sending syslog messages generally.

Default logging settings

Default logging settings displays the defaul logging preferences configured for the new access control rules. When enabled, new rules in all access control policies use the default settings automatically. You can override the settings for individual rules.

This section is view-only. To change the default logging preference, refer to *Access Control Preferences* in the *Cisco Secure Firewall Management Center Administration Guide*.

Default syslog settings

Default syslog settings provide two destination options:

- **Destination** > ***Specific Syslog alert object***: Events are sent based on the selected syslog alert as configured using the instructions in *Creating a Syslog alert Response* in the *Cisco Secure Firewall Management Center Administration Guide*. You can select the syslog alert from the list or add one by specifying the name, logging host, port, facility, and severity.

  When using this option, the system sends syslog messages to the server using the Management interface. Ensure there is a route from the Management interface to the syslog server, or messages will not arrive at the server.

- **Destination** > **Use FTD Platform Settings**: Connection or intrusion events are sent with the selected severity to syslog collectors configured in Platform Settings. Using this option, you can unify the syslog configuration by configuring it in platform settings and reusing the settings in access control policy. Severity selected in this section is applied to all connection and intrusion events. The default severity is alert.

Intrusion settings

Intrusion settings include these options:

- **Send syslog messages for intrusion events**: Send intrusion events as syslog messages. The defaults set above are used unless you override them.

- **Show overrides** or **Hide overrides**: If you want to use the default syslog destination and severity, leave these options empty. Otherwise, you can set a different syslog server destination for intrusion events, and change the severity of the events.

File and malware settings

File and malware settings include these options:

- **Send syslog messages for file and malware events**: Send file and malware events as syslog messages. The defaults set above are used unless you override them.
- **Show overrides** or **Hide overrides**: If you want to use the default syslog destination and severity, leave these options empty. Otherwise, you can set a different syslog server destination for file and malware events, and change the severity of the events.

Associate prefilter, decryption, and identity policies with an access control

Associate subpolicies with an access control policy to create a comprehensive security framework that handles different aspects of network traffic processing.

The easiest way to associate subpolicies to an access control policy is by clicking the policy's link in the packet flow shown at the top of an access control policy. Alternatively, you can use the policy's advanced settings to associate the subpolicy, as described in this topic. The subpolicies that you can associate with access control policy are:

- Prefilter policy—Performs early traffic handling using limited network (layer 4) outer-header criteria.
- Decryption policy—Monitors, decrypts, blocks, or allows application layer protocol traffic encrypted with Secure Socket Layer (SSL) or Transport Layer Security (TLS).
- Identity policy—Performs user identification based on the realm and authentication method associated with the traffic.

Before you begin

Before associating a decryption policy with an access control policy, review the information about TLS server identity discovery in TLS server identity discovery.

Procedure

Step 1 In the access control policy editor, select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 Locate the policy settings area for prefilter, decryption, or identity policies.

Step 3 Click **Edit** (✎) in the corresponding policy settings area.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 Choose a policy from the drop-down list.

If you choose a user-created policy, you can click the edit icon that appears to edit the policy.

Step 5 Click **OK**.

Step 6 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes.

Identify and fix anomalies with policy analyzer & optimizer

Use the Policy Analyzer & Optimizer to evaluate access control policies for anomalies such as redundant or shadowed rules, and take action to fix discovered anomalies.

The Policy Analyzer & Optimizer is hosted in the cloud and is different from the rule analysis available when you are not integrated with the cloud. The non-cloud policy analysis is not available once you integrate with the cloud.

The system automatically performs policy analysis on a daily basis (every 24 hours). You can also manually start an analysis. When you initially enable the service, the system starts an analysis of all existing access control policies.

Note Before optimizing a policy, create a copy of the policy. If you are then dissatisfied with the results of optimization, you can easily reassign the managed devices to the copy and return the system to its starting state.

Before you begin

- If you have enabled change management, Policy Analyzer & Optimizer automatically creates a ticket for the changes, and submits the ticket. The approver must approve the ticket before the changes can be deployed.
-
- Policy Analyzer & Optimizer adds rule comments on rules that are updated, disabled, or merged. You can later search on these comments to find optimized rules.
- Changes implemented by Policy Analyzer & Optimizer are reflected in the audit log as API calls under the default name internaladmin.
- If real-time Policy Analyzer & Optimizer is enabled, the system checks for applicable anomalies when you apply a new or modified rule.

Procedure

Step 1 Choose **Policies** > **Security policies** > **Access Control**.

If you have already run an analysis, the Anomaly column shows the number of issues with the policy and the percentage the policy can be optimized, and the state of the policy analysis, such as Error or Completed. Last Analyzed shows the date/time when the analysis was run.

Step 2 Select one or more policies, then click **Analyze Policy**.

The analysis runs as a background process in the cloud. When the analysis is complete, the results appear in the Anomaly column.

Notes:

- You can also start an analysis when editing the policy by selecting **Analyze** > **Policy**. Other options from that menu allow you to show hit counts and warnings.
- If you have not connected to the cloud yet, the explanatory dialog that opens when you click this button includes an **Integrate** button to help get you started. Policy Analyzer & Optimizer operates in the cloud only.

Step 3 When the analysis is complete, click the **% Optimizable** link in the Anomaly column to launch Policy Analyzer & Optimizer in the cloud.

When you have done all the actions you want to take, click **Apply Remediations** (in the cloud). You are shown a confirmation of what will be done. Click **Proceed** to implement the changes.

If the initial analysis ended in an error, you could instead click **Re-analyze** to restart the process.

Step 4 Deploy the policy to complete the changes.

If you have change management enabled, the approver must first approve the ticket that contains the remediations before you can deploy them.

View rule hit counts

Hit count indicates the number of times a policy rule or default action has been matched to a connection. The hit count is incremented only for the first packet of a connection that matches a rule. You can use this information to identify the efficacy of your rules. Hit count information is available only for access control and prefilter rules applied to Firewall Threat Defense devices.

Note

- The count persists through reboots and upgrades.
- Counts are maintained by each unit in an HA pair or cluster separately.
- You will not be able to derive the hit count information from a device when deployment or a task is in progress on the device.
- You can also see rule hit count information in the device CLI using the **show rule hits** command.
- If you have accessed the Hit Count page from the Access Control Policy page, you will not be able to view or edit prefilter rules and vice-versa.
- Hit counts are not available for rules that use the Monitor action.

Procedure

Step 1 In the access control policy or prefilter policy editor, click **Analyze** > **Manage Rule Hit Counts** on the top-right of the page.

Step 2 On the Hit Count page, select the device from the **Select a device** drop-down list.

If it is not the first time that you are generating hit counts for this device, the last fetched hit count information appears next to the drop-down box. Also, verify the **Last Deployed** time to confirm recent policy changes.

Select **All** to see an aggregate of hit counts across all devices assigned to the policy.

Step 3 If necessary, click **Refresh** ( ) to obtain current hit count data from the selected device.

In the prefilter policy, you might need to click **Fetch Current Hit Count** to get initial hit count data.

You cannot refresh the hit count while deployment to the device is in progress.

Step 4 View and analyze the data.

You can do these actions:

- Click **Prefilter** or **Access Control** to switch between the hit counts for these policies.
- Search for a specific rule by entering a search string in the **Filter** box.
- Broadly limit the list to **Hit Rules** or **Never Hit Rules** by selecting these options in the **Filter by** field. When viewing hit rules, you can further limit the list by selecting a time range in the **In Last** field (for example, in the last 1 day).
- (When viewed from the access control policy.) Clear the hit counts for one or more rule by selecting the checkbox for the rule, then clicking **Clear Hit Counts**. When confirming the action, choose **Clear and Reload** to refresh the hit count data. You can clear hit counts for up to 500 rules at a time. You cannot undo clearing hit counts.

Note Click the checkbox in the table header to select all rules visible in the list. To select a range of rules, select the first rule's checkbox, then **`Shift+click`** the last rule's checkbox; all rules in between are also selected.

- (When viewed from the access control policy.) You can do these actions with individual rules:
  - Edit the rule by clicking **Edit** from the **More** ( ) menu.
  - Delete the rule from the policy by clicking **Delete** from the **More** ( ) menu.
  - Enable or disable the rule by clicking **Enable Rule** or **Disable Rule** from the **More** ( ) menu.
  - Clear the hit count (reset it to zero) for the rule by clicking **Clear Hit Count** from the **More** ( ) menu. You cannot undo this action.
- (When viewed from the access control policy.) You can do these actions with individual rules:
  - Edit the rule by clicking **Edit** ( ).
  - Delete the rule from the policy by clicking **Delete** ( ).
  - Enable or disable the rule by clicking the **Slider** ( ).
  - Clear the hit count (reset it to zero) for the rule by clicking the **X** for the rule. You cannot undo this action.

- Change the displayed columns by clicking and selecting the columns to show. When viewed from the prefilter policy, use the **Cog** ( ) to select the columns to display.
- (When viewed from the prefilter policy.) Click on a rule name to edit it, or click **View** ( ) in the last column to view the rule details. Clicking on the rule name highlights it in the policy page where you can edit it.
- (When viewed from the prefilter policy.) Clear the hit count information (reset it to zero) for a rule by right-clicking the rule and selecting **Clear Hit Count**. You can select multiple rules by using Ctrl+click. You cannot undo this action.

Step 5 Optionally, you can generate a comma-separated values report of the details on the page by clicking **Generate CSV** on the bottom-left of the page.

Step 6 Click **Close** to return to the policy page.

Analyzing rule conflicts and warnings

View warnings and information about rule conflicts to examine the logic of your access control policy and to identify rules that need changes.

When rules overlap, you can end up with unnecessary rules in the policy, and these rules will never be matched to traffic. The analysis can help you delete unnecessary rules, or identify rules that should be moved or modified so they enforce the desired policy.

Policy warnings and errors point out things that you should understand and perhaps address to ensure that your rules provide the desired services.

Rule conflict analysis identifies these problems:

- Object overlap—An element in a field of a rule is a subset of one or more elements in the same field of the rule. For example, the source field might include a network object for 10.1.1.0/24, and another object for the host 10.1.1.1. Because 10.1.1.1 is within the network covered by 10.1.1.0/24, the object for 10.1.1.1 is redundant and can be deleted, simplifying the rule and saving device memory.
- Redundant rule—Two rules apply the same action to the same type of traffic and removing the base rule would not change the ultimate result. For example, if a rule permitting FTP traffic for a particular network were followed by a rule allowing IP traffic for that same network, and there were no rules in between denying access, then the first rule is redundant and you can delete it.
- Shadowed rule—This is the reverse of a redundant rule. In this case, one rule will match the same traffic as another rule such that the second rule will never be applied to any traffic because it comes later in the access list. If the action for both rules is the same, you can delete the shadowed rule. If the two rules specify different actions for traffic, you might need to move the shadowed rule or edit one of the two rules to implement your desired policy. For example, the base rule might deny IP traffic, and the shadowed rule might permit FTP traffic, for a given source or destination.

Before you begin

When doing the analysis:

- Only the first conflict for a given rule is identified. Once you fix the problem, the rule might be identified as having a conflict with another rule in the table. However, a rule might have multiple warnings or errors.

- Rule conflict analysis considers source/destination security zone, network, VLAN, and service/port match conditions and action only. It does not consider other match criteria, so an apparently redundant rule might not be completely redundant.
- FQDN network objects cannot be analyzed for conflict, because the IP address of an FQDN cannot be known prior to DNS lookup.
- Disabled rules are ignored.
- Time range attributes are ignored. Rules for different time periods might be marked as redundant when they are not actually redundant for the time ranges.
- Icons for warnings and errors, and rule conflicts when you enable the feature, are shown in the rules table. For a reference to the icons, see Rule and other policy warnings, on page 5.

Procedure

Step 1 Choose **Policies** > **Security policies** > **Access Control** and edit an access control policy.

Step 2 Do one of these to open the rule conflicts and warnings dialog box:

- To view rule conflicts, click the **Analyze** drop-down and click **Show Rule Conflicts** from the menu.

  You must re-enable rule conflict detection each time you open the policy or make a change and save the policy.
- To view rule warnings and errors, click **Analyze** > **Show Rule Warnings and Errors**.

  After you make a change to the policy, you can refresh the results by clicking **Re-Analyze** next to the **Analyze** button.
- To view policy warnings, click **Analyze** > **Show Policy Warnings**.

Step 3 In the rule conflicts and warnings dialog box:

- Rule warnings and errors are shown on a separate tab from rule conflicts. There is also a separate tab for policy warnings.
- Each tab contains sub-tabs to let you examine individual types of problems, such as redundant vs. shadowed or warnings vs. errors. You can also search for an item.
- **More** (⋮) next to each rule name provides shortcuts to edit, disable, or delete the rule.

Step 4 Click **Close** when finished.

Search for rules

Use search to help you find rules, especially when you have a large number of them.

Procedure

Step 1 When editing an access control policy, build the search string by clicking in the **Search** box.

- For a simple text string search, type the string. The search returns rules that have that string in any column. You cannot do a string search in conjunction with a tag search, such as combining a string search with a source networks search.
- To search on a specific column, start typing the name of the column until the system prompts you with the full name, such as Source Networks, or select the name from the list of searchable fields. When you select the search tag, you can then enter the search string for that tag. For example, **Source Networks 10.1.1.1**.
- When searching for IP addresses within the source/destination columns (but not as a simple text search), you can use the **Search Inside Subnets and Ranges** option to control whether you exactly match the IP address (unchecked) or match any subnet or range that includes the searched address (checked). For example, with the option checked, searching for 10.1.1.1 will include rules for 10.1.1.0/24.
- Searches in the port field return exact match only.
- When entering multiple values, separate the values with commas.
- When building a search string with multiple tags, do not include spaces between the tags.
- When you select a tag, you are prompted with values that appear in those columns. Select the values you want to search for.
- You can quickly filter based on some common features by clicking the filter icon to the left of the search box and selecting to show rules with any combination of the following: Allow, Block, Monitor, Intrusion Policy, Time Range, expired rules, rules that include objects whose definitions overlap..
- To see the rules that apply to a specific device or set of devices, click the filter icon and select the devices. Rules apply to a device if they use security zones that contain at least one interface on the device, or if they do not include security zones.

Step 2 With your cursor at the end of the search string in the search box, press Enter.

The rules that satisfy the search string are highlighted and non-matching rules are hidden. You can deselect **Show Only Matching Rules** to see the entire table, with the highlighted rules within the table. This lets you see the surrounding rules.

Beside the Show Only Matching Rules checkbox is a summary of the total number of rules in the policy compared to the number that match the search string.

Step 3 After your first search, clicking in the search box prompts you with recent searches and tags. You can quickly repeat a search by selecting it, or build similar searches by selecting previous searches or tags and building on them.

Step 4 To close the search and return to the unfiltered and unhighlighted table, click the **X** at the right of the search box. You can also put your cursor at the end of the search string and press the Esc key.